

## File Systems on Marc2

|                       |   |
|-----------------------|---|
| <code>/home</code>    | Network file system, accessible from all nodes, backed up regularly. Slow (compared to <code>/scratch</code> ), writing lots of IO or massively parallel IO to <code>/home</code> should be avoided.  |
| <code>/scratch</code> | Network file system, accessible from all nodes, no backup. Faster than <code>/home</code> , should be used for jobs that generate high amounts of IO.   |
| <code>/local</code>   | Local file systems on the compute node. Only accessible for processes on the same node. May be used for temporary files. Please remove your files or your job script from <code>/local</code> at the end of your program or when you finished your calculation. |

## Parameter limits

Marc2

|                     |  |
|---------------------|--|
| <code>h_rt</code>   | Mandatory, max. 10 days (864000s), no default.   |
| <code>h_vmem</code> | RAM per core, max. for 64 cores on single node: 4000M, default 1000M.  |
| <code>-pe</code>    | Mandatory for multicore, recommended: use <code>smp*</code> for multithreaded jobs (max. 64 cores) and <code>orte*</code> for multiprocess jobs. |

## qstat

Displaying queued jobs

|                               |  |
|-------------------------------|--|
| <code>-u &lt;user&gt;</code>  | Displays queued jobs for <code>&lt;user&gt;</code> . |
| <code>-u "*"</code>           | Display all queued jobs.                             |
| <code>-j &lt;jobID&gt;</code> | Display detailed info for a job.                     |
| <code>-j &lt;jobID&gt;</code> | Explain why this job is in error state.              |
| <code>-explain E</code>       |  |

## qalter

Changing script parameters

|   |   |
|---|---|
| <code>qalter [options] &lt;jobID&gt;</code> | Use to change or add any <code>qsub</code> parameters like <code>-pe</code> or <code>-l h_rt</code> . If a job is not allowed to start due to wrong parameters, it will run after setting them with <code>qalter</code> . |
| <code>qalter -w v &lt;jobID&gt;</code>      | Display queueing information for this job. Should display "verification: found suitable queue(s)," otherwise it is likely that the job's <code>qsub</code> parameters contain errors.                                     |

## Other useful commands

|                                     |  |
|-------------------------------------|--|
| <code>qdel &lt;jobID&gt;1</code>    | Delete a queued or running job.                          |
| <code>qmod -cj &lt;jobID&gt;</code> | Clear error state of a job.                              |
| <code>qmod -cj "*"</code>           | Clear error state for all of your jobs.                  |
| <code>dos2unix &lt;file&gt;</code>  | Convert a text file from DOS/Windows to UNIX formatting. |

## Module

Setting program environments

|   |  |
|---|--|
| Syntax  | <code>module &lt;command&gt;</code><br><code>&lt;modulename&gt;</code> |
| <code>avail</code>                                  | Display all available modules.   |
| <code>list</code>                                   | Display all loaded modules.  |
| <code>load &lt;module&gt;</code>                    | Load a module.   |
| <code>unload &lt;module&gt;</code>                  | Unload a module.   |
| <code>switch &lt;module1&gt; &lt;module2&gt;</code> | Unload <code>module1</code> , load <code>module2</code> .              |
| <code>purge</code>                                  | Unload all currently loaded modules.                                   |
| <code>. /etc/profile.d/modules.sh</code>            | Needed to use the <code>module</code> commands in job scripts.         |

## Possible job states in qstat

|                |         |
|----------------|---------|
| <code>q</code> | Queued  |
| <code>w</code> | Waiting |
| <code>h</code> | On hold |
| <code>r</code> | Running |
| <code>d</code> | Deleted |
| <code>e</code> | Error   |

## qsub

Submitting job scripts

|  |   |
|--|---|
| Syntax                                     | <code>qsub [options] jobscript [script options]</code>  |
| <code>-cwd</code>                          | Use current directory as working directory for a job.   |
| <code>-o &lt;path&gt;</code>               | Name and path to the output file.   |
| <code>-e &lt;path&gt;</code>               | Name and path to the error file.  |
| <code>-j</code>                            | Merge output and error streams into one file.   |
| <code>-</code>                             | resource-reservation,   |
| <code>l &lt;resource=value&gt;</code>      | e.g. <code>h_rt=&lt;max runtime&gt;</code> ,<br><code>h_vmem=&lt;max memory&gt;</code> .  |
| <code>-</code>                             | Choose the parallel environment <code>&lt;pe&gt;</code>   |
| <code>pe &lt;pe&gt; &lt;n_cores&gt;</code> | and <code>&lt;n_cores&gt;</code> to be used, e.g. <code>-pe orte* 16</code> for OpenMPI jobs with 16 cores. Omit this parameter for single core jobs.                                     |
| <code>-q &lt;queue&gt;</code>              | Specify the queue for job execution. Choose from <code>debug</code> , <code>regular</code> or <code>long</code> with different runtime limits (cf. Parameter limits, <code>h_rt</code> ). |
| <code>-V</code>                            | Export calling environment to job.  |
| <code>-N &lt;name&gt;</code>               | Set job name.   |
| <code>-t n[-m[:s]]</code>                  | Submit a task array with either amount of tasks <code>n</code> , index range <code>n-m</code> or range and step size <code>n-m:s</code> (ex. <code>-t 2-10:2</code> → 2,4,6,8,10).        |

## Contact, Further Information

HPC Hessen

[marc@hrz.uni-marburg.de](mailto:marc@hrz.uni-marburg.de) | [marburg@hpc-hessen.de](mailto:marburg@hpc-hessen.de)

## Detailed Information

Marc2

URL: <http://marc2-doc.hrz.uni-marburg.de/trac>

